

SEARCH-ON-THE-FLY WITH MERGE FUNCTION

Related Applications

This non-provisional application claims the benefit of U.S. provisional patent Application Number 60/227,305, entitled "SEARCH-ON-THE-FLY WITH MERGE FUNCTION," filed on August 24, 2000. The provisional application is hereby incorporated by reference.

This application is a continuation-in-part of Application Number 09/513,340, filed February 25, 2000, entitled Search-On-The-Fly/Sort-On-The-Fly Search Engine, which is hereby incorporated by reference.

Technical Field

The technical field is information management systems, interfaces, and mechanisms, and methods for searching one or more databases.

Background

In the most general sense, a database is a collection of data. Various architectures have been devised to organize data in a computerized database. Typically, a computerized database includes data stored in mass storage devices, such as tape drives, magnetic hard disk drives and optical drives. Three main database architectures are termed hierarchical, network and relational. A hierarchical database assigns different data types to different levels of the hierarchy. Links between data items on one level and data items on a different level are simple and direct. However, a single data item can appear multiple times in a hierarchical database and this creates data redundancy. To eliminate data redundancy, a network database stores data in nodes having direct access to any other node in the database. There is no need to duplicate data since all nodes are universally accessible. In a relational database, the basic unit of data is a relation. A relation corresponds to a table having rows, with each row called a tuple, and columns, with each column called an attribute. From a practical standpoint, rows represent records of related data and columns identify individual data elements. The order in which the rows and columns appear in a table has no significance. In a relational database, one can add a new column to a table without having to modify older applications that access other columns in the table. Relational databases thus provide flexibility to accommodate changing needs.

All databases require a consistent structure, termed a schema, to organize and manage the information. In a relational database, the schema is a collection of tables. Similarly, for each table, there is generally one schema to which it belongs. Once the

1 schema is designed, a tool, known as a database management system (DBMS), is used to
2 build the database and to operate on data within the database. The DBMS stores,
3 retrieves and modifies data associated with the database. Lastly, to the extent possible,
4 the DBMS protects data from corruption and unauthorized access.

5 A human user controls the DBMS by providing a sequence of commands selected
6 from a data sublanguage. The syntax of data sublanguages varies widely. The American
7 National Standards Institute (ANSI) and the International Organization for
8 Standardization (ISO) have adopted Structured English Query Language (SQL) as a
9 standard data sublanguage for relational databases. SQL comprises a data definition
10 language (DDL), a data manipulation language (DML), and a data control language
11 (DCL). The DDL allows users to define a database, to modify its structure and to destroy
12 it. The DML provides the tools to enter, modify and extract data from the database. The
13 DCL provides tools to protect data from corruption and unauthorized access. Although
14 SQL is standardized, most implementations of the ANSI standard have subtle differences.
15 Nonetheless, the standardization of SQL has greatly increased the utility of relational
16 databases for many applications.

17 Although access to relational databases is facilitated by standard data
18 sublanguages, users still must have detailed knowledge of the schema to obtain needed
19 information from a database since one can design many different schemas to represent the
20 storage of a given collection of information. For example, in an electronic commerce
21 system, product information, such as product SKU, product name, product description,
22 price, and tax code, may be stored in a single table within a relational database. In
23 another electronic commerce system, product SKU, product name, description, and tax
24 code may be stored in one table while product SKU and product price are stored in a
25 separate table. In this situation, a SQL query designed to retrieve a product price from a
26 database of the first electronic commerce system is not useful for retrieving the price for
27 the same product in the other electronic system's database because the differences in
28 schemas require the use of different SQL queries to retrieve product price. As a
29 consequence, developers of retail applications accessing product information from
30 relational databases may have to adapt their SQL queries to each individual schema.
31 This, in turn, prevents their applications from being used in environments where there are
32 a wide variety of databases having different schemas, such as the World Wide Web.

33 A further problem with conventional searches, search engines, data access and
34 data retrieval is a tendency to return very large amounts of data, or to require the search

1 menus are created "on-the-fly" as the user steps through the sort and/or search process.
2 Thus, search-on-the-fly is inherently intuitive, and allows a user with little or no
3 knowledge of the database contents, its organization, or a search engine search routine to
4 execute comprehensive analysis, sorting and/or searches that return generally accurate
5 results.

6 Search-on-the-fly also searches on key words specified by the user. Search-on-
7 the-fly can be used to exclude certain items. Search-on-the-fly incorporates other
8 advanced features such as saving results by attaching a cookie to a user's computer, and
9 associating icons with the results.

10 Search-on-the-fly may be used with both internal and external databases. For
11 example, Search-on-the-fly may be used with a company internal database and one or
12 more databases accessible through the Internet.

13 Search-on-the-fly is user-friendly. With one interface, many different types of
14 databases or database schemas may be searched or sorted.

15 Finally, the search-on-the-fly technique, and other techniques discussed above
16 may be used in conjunction with a method of doing business, particularly a business
17 method that uses the Internet as a communications backbone.

18 **Description of the Drawings**

19 The detailed description will refer to the following figures, in which like numerals
20 refer to like objects, and in which:

21 Figure 1 is a block diagram of a system that uses a search-on-the-fly/sort-on-the-
22 fly process;

23 Figure 2 is another overall block diagram of the system of Figure 1;

24 Figure 3 is a detailed block diagram of the search engine used with the system of
25 Figure 2;

26 Figure 4 is an example of a search-on-the-fly using the search engine of Figure 3;

27 Figures 5 - 9 are detailed block diagrams of components of the search engine of
28 Figure 3;

29 Figure 10 is another example of a search-on-the-fly using the search engine of
30 Figure 3;

31 Figures 11 - 15b are additional examples of a search-on-the-fly using the search
32 engine of Figure 3;

33 Figures 16 - 20 are flow charts illustrating operations of the search engine of
34 Figure 3;

1 Figure 21 illustrates a further function of the search engine of Figure 3 in which
2 results of more than one search are combined;

3 Figures 22 - 26 illustrate graphical user interfaces that may be displayed in
4 conjunction with operation of the system of Figure 1;

5 Figure 27 is a flowchart illustrating an alternate operation of a query generator
6 used with the search engine of Figure 3;

7 Figure 28 is a flowchart illustrating an alternate operation of the truncator used
8 with the search engine of Figure 3;

9 Figures 29 - 36 illustrate user interfaces with search results from a search on the
10 fly and a merge function;

11 Figures 37 - 39 illustrate a keyword search result form a search on the fly with the
12 merge function;

13 Figures 40-49 illustrate additional search results;

14 Figure 50 illustrates a cellular phone incorporating the search-on-the fly with
15 merge function;

16 Figure 51 illustrates a personal data assistant incorporating the search-on-
17 the-fly with merge function;

18 Figures 52a - 52l illustrate search-on-the-fly as displayed on the cellular phone of
19 Figure 50; and

20 Figure 53 illustrates a computer-readable medium having the search-on-the-fly
21 with merge function loaded thereon.

22 **Detailed Description**

23 Ordinary search engines place constraints on any search. In particular, a partial
24 ordering of available search criteria limits application of the search engine only to certain
25 search sequences. The user is given a choice of search sequences, and the order in which
26 individual search steps in the search sequence become available limits the direction of the
27 search. A user who desires to take a vacation cruise may use an Internet search engine to
28 find a desired vacation package. The search begins with presentation of a list of general
29 categories, and the user clicks on "travel," which produces a list of subcategories. The
30 user then clicks on "cruises" from the resulting list of subcategories, and so on in a
31 cumulative narrowing of possibilities until the user finds the desired destination, date,
32 cruise line, and price. The order in which choices become available amounts to a
33 predefined "search tree," and the unspoken assumption of the search engine designer is

1 that the needs and thought processes of any user will naturally conform to this predefined
2 search tree.

3 To an extent, predefined constraints are helpful in that predefined constraints
4 allow a search engine to logically and impersonally order the user's thoughts in such a
5 way that if the user has a clear idea of what object the user wants, and if the object is
6 there to be found, then the user is assured of finding the object. Indeed, the user may
7 want to know that choosing any available category in a search sequence will produce an
8 exhaustive and disjunctive list of subcategories from which another choice can be made.
9 Unfortunately, an unnecessarily high cost is too often paid for this knowledge: The user
10 is unnecessarily locked into a limited set of choice sequences, and without sufficient prior
11 knowledge of the object being sought, this limitation can become a hindrance.
12 Specifically, where prescribed search constraints are incompatible with the associative
13 relationships in the user's mind, a conflict can arise between the thought processes of the
14 user and the function of the search engine.

15 At one time, such conflicts were written off to the unavoidable differences
16 between computers and the human mind. However, some "differences" are neither
17 unavoidable nor problematic. In the case of search engine design, the solution is elegant:
18 upon selecting a category or entering a keyword, the user can be given not only a list of
19 subcategories, but the option to apply previously available categories as well. In slightly
20 more technical terms, the open topology of the search tree can be arbitrarily closed by
21 permitting search sequences to loop and converge. Previous lists can be accessed and
22 used as points of divergence from which new sub-sequences branch off, and the attributes
23 corresponding to distinct sub-sequences can later be merged.

24 Sort-on-the-fly/search-on-the-fly data analysis, sorting access and retrieval
25 methods and apparatus (hereafter, search-on-the-fly search engine) provide an intuitive
26 means for analyzing various types of databases, allowing a user to obtain information
27 about and/or access data in the database without having to know anything about the
28 database structure. A user selects a desired term, and a database manager reviews the
29 database for all instances of the desired term, even if a specific file or table does not
30 contain the instance. For example, if a user wants to analyze the database using the name
31 of a specific individual as a database entry point, the database manager will search the
32 database or index using the desired name, and will organize the results so that all entries
33 associated with that name are displayed. The database need not have a specific file (in a
34 flat database) or a table (in a relational database) of names. The user may perform further

1 on-the-fly searches to narrow the search results, or for other reasons. The search engine
2 then conducts a further search using this criteria and produces a second search result.
3 Further narrowing or broadening of the search are permitted, with the search engine
4 returning results based on any new criteria.

5 This on-the-fly method or process can be used to simply analyze data or gather
6 information about data stored in a database. The actual data itself does not need to be
7 fetched, displayed, printed or even sorted. The user may simply wish to use this tool to
8 "clean-up" data or understand how data could be sorted or for other reasons.

9 Figure 1 is a block diagram of a system 10 that uses search-on-the-fly. In Figure
10 1, a database 12 is accessed using a hardware/software interface device 100 to provide
11 data to a user terminal 14. Additional databases 13 and 15 may also be accessed by the
12 terminal 14 using the device 100. The databases 12, 13 and 15 may use different
13 schemas, or may use a same schema. As will be described later, the device 100 may
14 include the search-on-the-fly search apparatus. In an alternative embodiment, the search-
15 on-the-fly search engine may be co-located with the terminal 14. In yet another
16 embodiment, the search-on-the-fly search engine may be incorporated into the structure
17 of one or more of the databases 12, 13 and 15. The device 100 may interface with any
18 one or more of the databases 12, 13 and 15 using a network connection such as through
19 the Internet, for example. Other communications mediums may also be used between the
20 terminal 14, the device 100 and any one or more of the databases 12, 13 and 15. These
21 mediums may include the public switched telephone network (PSTN), cable television
22 delivery networks, Integrated Services Digital Networks (ISDN), digital subscriber lines
23 (DSL), wireless means, including microwave and radio communications networks,
24 satellite distribution networks, and any other medium capable of carrying digital data.

25 The system shown in Figure 1 is but one of many possible variations. The search-
26 on-the-fly search engine could also be incorporated within a single computer, such as a
27 personal computer, a computer network with a host server and one or more user stations,
28 an intranet, and an Internet-based system, as shown in Figure 2. Referring again to Figure
29 2, the terminal 14 may be any device capable of displaying digital data including
30 handheld devices, cellular phones, geosynchronous positioning satellite (GPS) devices,
31 wrist-worn devices, interactive phone devices, household appliances, televisions,
32 television set top boxes, handheld computers, and other computers.

33 Figure 3 is a detailed block diagram of an exemplary search-on-the-fly search
34 engine 125. The search engine 125 includes a request analyzer 130 that receives search

1 The search engine 125 may conduct a search using one or more search cycles. A
2 search cycle includes receipt of a request 114, any necessary formatting of the request
3 114, and any necessary truncation steps. The search cycle ends when a result list 175 is
4 provided to the terminal 14. The search engine 125 may retain a status of each past and
5 current search cycle so that the user can modify the search at a later time. The user may
6 also use this feature of retaining a status of past and current search cycles to combine
7 results of multiple searches, using, for example, a Boolean AND function, a Boolean OR
8 function, or other logic function. The above listed functions will be described in more
9 detail later.

10 The search-on-the-fly function of the search engine 125 begins by determining
11 available data fields of the database 12. The database 12 may have its data organized in
12 one or more data fields, tables, or other structures, and each such data field may be
13 identified by a data field descriptor. In many cases, the data field descriptor includes
14 enough text for the user at the terminal 14 to determine the general contents of the data
15 field. The list of data fields may then be presented at the terminal 14, for example, in a
16 pull down list. An example of such a data field result list is shown in Figure 4, which is
17 from a federal database showing data related to managed health care organizations. This
18 database is available at <http://tobaccopapers.org/dnld.htm>. In Figure 4, the first data field
19 listed is "PlanType," which is shown in result list 156. Other data field descriptors show
20 the general categories of data in the database.

21 Using the terminal 14, the user may select one of the data field descriptors to be
22 searched. For example, the user could select "city." If a number of entries, or records, in
23 the city data field is short, a further result list of complete city names may be displayed.
24 If the entries are too numerous to be displayed within a standard screen size, for example,
25 the search engine 125 may, in an iterative fashion, attempt to reduce, or truncate, the
26 result list until the result list may be displayed. In the example shown in Figure 4, entries
27 in the city data field are so numerous (the database includes all U.S. cities that have a
28 managed health care organization) that the search engine 125 has produced a result list
29 157 that shows only a first letter of the city. Based on the available database data fields,
30 the user may then perform a further search-on-the-fly. In this case, the user may choose
31 cities whose first initial is "N." The search engine 125 then returns a result list 158 of
32 cities whose names start with the letter "N." Because in this instance the result list 158 is
33 short, no further truncation is necessary to produce a manageable list.

1 Figure 5 is a more detailed block diagram of the request analyzer 130. A protocol
2 analyzer 133 receives the request 114 and provides an output 135 to a constraint collator
3 136. The protocol analyzer 133 examines the received request 114, determines a format
4 of the request 114, and performs any necessary translations to make the request format
5 compatible with the database to be accessed. If the database to be accessed by the
6 terminal 14 is part of a same computer system as the terminal 14, then the protocol
7 analyzer 133 may not be required to perform any translations or to reformat the request
8 114. If the database to be accessed is not part of the same computer system as the
9 terminal 14, then the protocol analyzer 133 may be required to reformat the request 114.
10 The reformatting may be needed, for example, when a request 114 is transmitted over a
11 network, such as the Internet, to a database coupled to the network.

12 The constraint collator 136 provides the updated request 115 (which may be an
13 initial request, or a subsequent request) to the query generator 150. The constraint
14 collator 136 is responsible for interpreting the request 114. The constraint collator 136
15 performs this function by comparing the request 114 against information stored in the
16 status control 140. In particular, the constraint collator 136 sends the request status
17 control signal 118 to the status control 140 and receives the request status response 119.
18 The constraint collator 136 then compares the request status response 119 to constraint
19 information provided with the request 114 to determine if the constraint status should be
20 updated (e.g., because the request 114 includes a new constraint). In an embodiment, the
21 constraint collator 136 compares constraint information in a current request 114 to
22 constraint information residing in the status control 140, and if the current request 114
23 includes a new constraint, such as a new narrowing request (for example, when the user
24 clicks, touches or points over a field shown in a last search cycle), then the constraint
25 collator 136 adds the updated information and sends the updated request 115 to the query
26 generator 150. If the constraint status should be updated, the constraint collator 136 sends
27 the status update 118 to the status control 140. If the request 114 is a refresh request, the
28 constraint collator 136 sends a reset command 131 to the database qualifier 160. The
29 updated request 115 (possibly with a new constraint) is then sent to the query analyzer
30 150 for further processing.

31 Figure 6 is a block diagram of the query generator 150. The overall functions of
32 the query generator 150 are to scan a database, such as the database 12, using the database
33 driver 170, and to collect search results based on constraints supplied by the request

1 analyzer 130. The query generator 150 then returns the search results 175 to the terminal
2 14.

3 The query generator 150 includes a truncator 152 and a dispatcher 154. The
4 truncator 152 receives the updated request 115, including a new constraint, if applicable.
5 The truncator 152 creates new queries, based on new constraints, and applies the new
6 requests 151 to the database 12 using the database driver 170. Many different methods of
7 truncating for display or viewing may be used by truncator 152. The truncator 152 may
8 include a variable limit 155 that is set, for example, according to a capacity of the
9 terminal 14 to display the search results 175. If data retrieved from the database 12
10 exceed the limit value, the truncator 152 adjusts a size (e.g., a number of entries or
11 records) of the data until a displayable result list is achieved. One method of adjusting
12 the size is by cycling (looping). Other methods may also be used to adjust the size of the
13 result list. For example, the terminal 14 may be limited to displaying 20 lines of data
14 (entries, records) from the database 12. The truncator 152 will cycle until the displayed
15 result list is at most 20 lines. In an embodiment, the truncation process used by the
16 truncator 152 assumes that if the user requests all values in a particular data field from the
17 database 12, and there are no other constraints provided with the request 114, and if the
18 size of the resulting result list is larger than some numeric parameter related to a display
19 size of the terminal 14, then the constraints may be modified by the truncator 152 so that
20 the result list can accommodated (e.g., displayed on one page) by the terminal 14. For
21 example, instead of a full name of a city, some part of the name - the first n letters - is
22 checked against the database 12 again, and n is reduced until the result list is small
23 enough for the capacity of the terminal 14. If the maximum number of displayable results
24 is three (3), and the database 12 contains the names of six cities "Armandia, Armonk,
25 New Orleans, New York, Riverhead, Riverdale," then the first attempt to "resolve" the
26 result list will stop after a result list display is created with the full name of the cities:

27 Armandia, Armonk, New Orleans... (the limit was reached)

28 Try again with 7 characters:

29 Armandia, Armonk, New Orl, New Yor, (limit reached again)

30 Again with 5 characters:

31 Armandia, Armonk, New O, New Y, (limit reached again)

32 Again with 3 characters:

33 Arm (...), New (...), Riv (...). These results may now be displayed on the terminal 14.

34 The display of Arm, New, Riv can then be used to conduct a further search-on-the-fly.

1 For example, a user could then select Riv for a further search-on-the-fly. The result list
2 returned would then list two cities, namely Riverhead and Riverdale.

3 In another embodiment, a fixed format is imposed such that all queries generated
4 against a database will have preset limits corresponding to the capacity of the terminal 14.

5 In yet another embodiment, the truncator 152 may adjust the field size by division
6 or other means. For example, if the display limit has been reached, the truncator 125 may
7 reduce the field size, X by a specified amount. In an embodiment, X may be divided by
8 two. Alternatively, X may be multiplied by a number less than 1, such as $3/4$, for
9 example. Adjusting the field size allows the search engine 125 to perform more focused
10 searches and provides more accurate search results.

11 In another embodiment, the truncator first attempts to display information without
12 truncation. If that is not appropriate, the truncator may attempt truncation by beginning
13 with one character (26 letters and perhaps 10 digits) and incrementing to two characters
14 and then three, four, until a failure to display is reached.

15 In still another embodiment, the user may select a limit that will cause the
16 truncator 152 to adjust the field size. For example, the user could specify that a
17 maximum of ten entries should be displayed.

18 For certain data fields, a terminal of a hand-held device, may have a very limited
19 display capacity. For example, a personal data assistant (POA - see Figure 52) or a
20 cellular phone (see Figure 50) may be used to search a database, with the results
21 displayed on a small screen. Alternatively a user may specify a limit on the number of
22 entries for display. In the illustrated cases, the search engine 125 may return a result list
23 175 of the request 114 on multiple display pages, and the user may toggle between these
24 multiple display pages. As an example, if the terminal 14 is limited to displaying a
25 maximum of ten entries, and if the request 114 results in a return of a data field
26 comprising the 400 largest cities in the United States, the truncator 152 will produce a list
27 of 23 entries comprising 23 alphabetical characters (no cities that begin with Q, Y or Z -
28 see Figure 4). The search engine 125 may then display the results on three pages.
29 Alternatively, the truncator 152 could produce a list of letter groups into which the cities
30 would fall, such as A-D, E-G, H-M, N-R, and R-X, for example. In another alternative,
31 the search engine 125 may send a notice to the terminal that the request 114 cannot be
32 accommodated on the terminal 14 and may prompt the user to add an additional
33 constraint to the request 114, so that a search result may be displayed at the terminal 14.

1 Adjusting the data field size also provides more convenient search results for the
 2 user. For example, if a user were to access an Internet-based database for books for sale,
 3 and were to request a list of all book titles beginning with the letter "F," a common search
 4 engine might return several hundred titles or more, displaying perhaps twenty titles
 5 (entries) at a time. The user would then have to look through each of many pages to find
 6 a desired title. This process could be very time-consuming and expensive. Furthermore,
 7 if the search results were too large, the common search engine might return a notice
 8 saying the results were too large for display and might prompt the user to select an
 9 alternative search request. However, performing the same search using the search engine
 10 125 allows the truncator 152 to reduce the size of the information displayed to a
 11 manageable level. In this example, if the request 114 includes the constraint "F," the
 12 truncator 152 will loop through the data in a data field that includes book titles starting
 13 with the letter "F" until a list is available that can fit within the display limits of the
 14 terminal 14, or that fits within a limit set by the user, for example. The first list returned
 15 to the terminal 14 as a result of this request 114 may be a two letter combination with "F"
 16 as the first letter and a second letter of a book title as the second letter. For example, the
 17 first list may include the entries "Fa," "Fe," "Fi," "Fo," and "Fu," all of which represent
 18 titles of books. The user could then select one of the entries "Fa," "Fe," "Fi," "Fo," and
 19 "Fu" to perform a further search, continuing the process until one or more desired titles
 20 are displayed. An example of a similar truncation result is shown in Figure 14.

21 When a parameter related to the search results is adequately truncated, the
 22 parameter is directed to the dispatcher 154, which retrieves the data from database 12
 23 using the database driver 170. The dispatcher 154 then directs the final, truncated search
 24 results 175 back to the terminal 14 as a response to the request 114.

25 Figure 7 is a block diagram showing the status control 140, which is responsible
 26 for monitoring the status of a current search. Due to the nature of the search engine 125,
 27 the user can choose any combination of constraints, fields or keywords, including those
 28 from past and current search cycles. The status control 140 may keep track of all past
 29 cycles of the search, as well as all information necessary to return to any of those past
 30 search cycles. The status control 140 includes a status data module 142, and an index
 31 module 144. The status data module 142 contains data related to each such search cycle,
 32 including the constraint(s) entered during the search cycle, any truncation steps taken, and
 33 the results of such truncation, for example. The index module 144 provides access to
 34 these data. When the request 114 is being analyzed by the request analyzer 130, the

1 constraint collator 136 sends a request status query 116 to the index module 144. The
2 status data module 142 contains information related to all past and current search cycles,
3 which are referenced by the index module 144, and delivers a status response 119 for the
4 most recent search cycle to the constraint collator 136. When a new constraint is sent to
5 the query generator 150, the status data module 142 is updated 118 by the constraint
6 collator 136. Specific structures of the request 114, the request status query 116, the
7 status response 119 and the request status control 118 will be provided later.

8 The status data module 142 may be reset by the database qualifier 160 with all
9 available fields when a refresh function is used. In an embodiment, the refresh function
10 may be used to clear all past search cycles and the current search cycle from the status
11 control 140. In such an event, the search results, such as the search results shown in
12 Figure 4, will no longer be displayed at the terminal 14, and data related to the past and
13 the current search cycles may not be used for future search cycles. In effect, the refresh
14 function may cause the entire search to be discarded. The refresh function may be
15 activated when a user selects a refresh button (see Figure 4) on a displayed result list, or
16 on another portion of a GUI. Alternatively, the refresh function may discard selected
17 search cycles. In this alternative embodiment, the user may, for example, move a cursor
18 to a desired result list from a past search cycle and activate a refresh, reset, back, or drop
19 button. All data associated with search cycles subsequent to the selected search cycle,
20 including all displayed result lists may then be discarded.

21 Figure 8 is a block diagram showing the database qualifier 160. The database
22 qualifier 160 provides data field information at the start of a search or when the search
23 engine 125 is refreshed. A field assessor 162 access the database 12 using the database
24 driver 170, and identifies and accesses discrete data fields and other information in the
25 database 12. A field converter 164 structures the data field information into a usable
26 (searchable/sortable) structure and sends 163 the formatted data field information to the
27 status control 140. Techniques for identifying and accessing the data fields, and for
28 formatting the data field information are well known in the art. Such techniques are
29 described, for example, in U.S. Patent 5,222,066, Interface for Accessing Multiple
30 Records Stored in Different File System Formats, the disclosure of which is hereby
31 incorporated by reference.

32 Figure 9 is a block diagram of the database driver 170. The database driver 170
33 is the universal interface with the database 12, which can be a local or a remote database.

1 Figure 10 is an example of a search-on-the-fly using the search engine 125. In
2 Figure 10, a database 200 includes information related to a number of individuals. The
3 information in the database 200 may be presented at the terminal 14 using a series of
4 screens or menus 201 - 230. The user first accesses the database 200 and is presented
5 with a list 201 of the information or data fields contained in the database 200. The result
6 list 201 is generated by the field assessor 162, and is provided for display at the terminal
7 14 by the query generator 150. As shown in Figure 10, a user has selected the data field
8 "City" for display of information. However, the number of "cities" listed in the database
9 200 is too large to conveniently display at one time (i.e., on one page) at the terminal 14.
10 Accordingly, the truncator 152 will loop a required number of times until an adequate
11 display is available. In Figure 10, the menu 203 shows the results of the truncation with
12 only the first letter of a city name displayed.

13 Using the menu 203, the user has selected cities beginning with the letter "A."
14 The results are shown in menu 205. Now, the user elects to conduct another search-on-
15 the-fly, by selecting the "sort-on-the-fly" option 206. The query generator 150 displays
16 all the information fields available from the database 200, except for the information field
17 already displayed, namely "City." The results are displayed in menu 207. The user then
18 elects to further search on the data field "State." The query generator 150 returns the
19 requested information as displayed in menu 209, listing five states by their common two-
20 letter abbreviation. The user then chooses New York from the menu 209, and the query
21 generator 150 returns a list of cities in New York, menu 211.

22 Next, the user elects to conduct another search-on-the-fly, option 212, and the
23 query generator 150 returns only the remaining data fields for display in menu 215. From
24 the menu 215, the user selects "Address" for the next data field to search, and the query
25 generator 150 returns an menu 217 showing only first letters of the address. This
26 signifies that the data field "Address" was too large to be easily displayed on the terminal
27 14. The user then elects to search on all addresses that begin with "C." The query
28 generator 150 returns a list of addresses by displaying only street names, menu 219.

29 The user then elects to conduct a further search-on-the-fly, option 220, and the
30 remaining two data fields, "Name" and "Phone" are displayed as options in menu 221.
31 The user selects name, and the query generator returns a further breakdown of the data by
32 last name and by first name, menu 223. This process continues, with further menus being
33 used to select a last name and a first name from the database 200. When the final

1 selection is made, information from the database 200 related to the individual is displayed
2 in window 230.

3 In the example shown in Figure 10, the user could have refreshed the search
4 engine 125 at any time, and the search would have recommenced at the beginning.
5 Alternatively, the user could, by simply selecting a prior menu, such as the menu 215,
6 have changed the course of the search. In this alternative, if the user had gone back to the
7 menu 215 and instead of selecting "Address" selected "Phone," then the menus 217 - 229
8 would be removed from display at the terminal 14, and the search would begin over from
9 the point of the menu 215.

10 Figures 11 – 15b illustrate exemplary searches of a remote database, such as the
11 database 13 shown in Figure 1. The database in the illustrated example is for an Internet
12 website 232 that sells books. The examples illustrated are based on a Barnes & Noble™
13 website. In Figure 11, the user has applied the search engine 125 to the website 232
14 database, and the query generator 150 has returned a list 233 of data fields from which the
15 user may select to access data from the website 232 database. The list 233, and other lists
16 described below, may be displayed as overlays on the website 232. In the example
17 illustrated, the user selects "Title" for the first search cycle. Because the list of titles is
18 too large to easily display at the terminal 14, the truncator 152 loops until an
19 alphanumeric list 234 is created. The list 234 is then returned to the terminal 14. For the
20 next search cycle, the user selects titles that begin with the letter "C." Again, the data
21 field contains too many entries to conveniently display at the terminal 14, and the
22 truncator 152 loops as appropriate until list 235 is created. The process continues with
23 subsequent lists 236 and 237 being returned to the terminal 14.

24 Figures 12 - 15b illustrate alternate searches that may be completed using the
25 website 232 database.

26 For the search results shown in Figures 11 – 15b, the status control 140 may
27 iterate as follows:

28 Status Control Started...

29 Key: Title1 Option: Title Level: 1 Filter: Field: Title

30 Key: A2 Option: A Level: 2 Filter: SUBSTRING([Title],1,1) = 'A' Field:

31 Title

32 Key: AA3 Option: AA Level: 3 Filter: SUBSTRING([Title],1,2) = 'AA'

33 AND SUBSTRING([Title],1,1) = 'A' Field: Title

1 Key: F4 Option: F Level: 4 Filter: SUBSTRING([Title],1,1) = 'F' Field:
2 Title
3 Key: Fa5 Option: Fa Level: 5 Filter: SUBSTRING([Title],1,2) = 'Fa'
4 AND SUBSTRING([Title],1,1) = 'F' Field: Title
5 Key: Favo6 Option: Favo Level: 6 Filter: SUBSTRING([Title],1,4) =
6 'Favo' AND SUBSTRING([Title],1,2) = 'Fa' AND SUBSTRING([Title],1,1) = 'F'
7 Field: Title
8 Key: C7 Option: C Level: 7 Filter: SUBSTRING([Title],1,1) = 'C' Field:
9 Title
10 Key: Ce8 Option: Ce Level: 8 Filter: SUBSTRING([Title],1,2) = 'Ce'
11 AND SUBSTRING([Title],1,1) = 'C' Field: Title
12 Key: Cells9 Option: Cells Level: 9 Filter: SUBSTRING([Title],1,5) =
13 'Cells' AND SUBSTRING([Title],1,2) = 'Ce' AND SUBSTRING([Title],1,1) = 'C'
14 Field: Title
15 Key: Cellula10 Option: Cellula Level: 10 Filter: SUBSTRING([Title],1,7)
16 = 'Cellula' AND SUBSTRING([Title],1,2) = 'Ce' AND SUBSTRING([Title],1,1)
17 = 'C' Field: Title
18 Key: CC11 Option: CC Level: 11 Filter: SUBSTRING([Title],1,2) = 'CC'
19 AND SUBSTRING([Title],1,1) = 'C' Field: Title
20 Status Control Terminated.
21 Figure 15b shows the results for a search for a low-fat cookbook using the search
22 engine 125 as applied to a remote database. In this example, the remote database is
23 coupled to a Barnes & Noble web page. The first query, and resulting message strings,
24 are illustrated by the following:
25 Query Analyzer
26 Message Received: ACK
27 Status Control: Refresh
28 Dispatcher
29 Message Sent: Categories~~Title~~Author~~ISBN~SubTitle~Format~Date
30 Published~Stock Status~Recommended
31 Age~Pages~Ratings~Price~Retail~Savings~~Publisher
32 Query Analyzer
33 Message Received: CLK#0#1#Categories
34 Status Control received an update:

1 Key: Categories1 Option: Categories Level: 1 Filter: Field: Categories
2 Query Generator
3 Request is not cached, processing
4 Generated Query: SELECT DISTINCT [Categories] FROM Books ORDER BY
5 [Categories]
6 Number of Matching Records: 2032
7 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,82) FROM Books
8 ORDER BY SUBSTRING([Categories],1,82)
9 Number of Matching Records: 2022
10 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,61) FROM Books
11 ORDER BY SUBSTRING([Categories],1,61)
12 Number of Matching Records: 1995
13 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,45) FROM Books
14 ORDER BY SUBSTRING([Categories],1,45)
15 Number of Matching Records: 1751
16 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,33) FROM Books
17 ORDER BY SUBSTRING([Categories],1,33)
18 Number of Matching Records: 1251
19 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,24) FROM Books
20 ORDER BY SUBSTRING([Categories],1,24)
21 Number of Matching Records: 799
22 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,18) FROM Books
23 ORDER BY SUBSTRING([Categories],1,18)
24 Number of Matching Records: 425
25 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,13) FROM Books
26 ORDER BY SUBSTRING([Categories],1,13)
27 Number of Matching Records: 319
28 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,9) FROM Books
29 ORDER BY SUBSTRING([Categories],1,9)
30 Number of Matching Records: 147
31 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,8) FROM Books
32 ORDER BY SUBSTRING([Categories],1,8)
33 Number of Matching Records: 111

1 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,7) FROM Books
2 ORDER BY SUBSTRING([Categories],1,7)
3 Number of Matching Records: 78
4 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,6) FROM Books
5 ORDER BY SUBSTRING([Categories],1,6)
6 Number of Matching Records: 44
7 Generated Query: SELECT DISTINCT SUBSTRING([Categories],1,5) FROM Books
8 ORDER BY SUBSTRING([Categories],1,5)
9 Number of Matching Records: 26
10 Truncator finished, took 15 seconds to make 13 iterations
11 Caching this request...
12 Dispatcher
13 Message Sent: Afric~Art,
14 ~Biogr~Busin~Compu~Cooki~Engin~Enter~Ficti~Histo~Home ~Horro~Kids!~Law:
15 ~Medic~Mind,~Nonfi~Paren~Poetr~Refer~Relig~Scien~Small~Sport~Trave~Write~
16 Query Analyzer
17 Message Received: CLKCategories

18 In the example illustrated by Figure 15b and the above-listed message strings, an
19 initial request would have returned 2032 book titles for cook books. This number of
20 entries may be too large. Accordingly, the truncator 152, through 13 iterations, reduces
21 the entries in a result list to 26. The entries in the truncated result list can then be easily
22 reviewed by the user, and further searches may be performed to identify a desired book.
23 As can be seen in Figure 15b, the user has selected "Categories" as a data field to search.
24 As is also shown in Figure 15b, the search engine 125 may display other information
25 windows, such as book availability, ordering and shipping information windows. With a
26 simple drag-and-drop cursor operation, for example, the user may then order and pay for
27 the desired book.

28 Figure 16 - 20 are flow charts illustrating operations of the search engine 125.
29 Figure 16 is a flowchart of an overall search routine 250. The process starts in block 251.
30 The request analyzer 130 receives the request 114, block 252. The request 114 may be
31 made using a hierarchical menu-based display or a graphical user interface, with one or
32 more layers. Using either the menu or the GUI, the user may enter specific details by
33 typing, selection of iconic symbols or pre-formatted text, and by using well-known data
34 entry techniques, for example. The request 114 may also comprise a simple text or voice

1 query. Use of voice recognition may be particularly useful in mobile environments, and
 2 to speed access to the database 12. Use of voice recognition may include simple
 3 commands, such as UP, DOWN, and SELECT, to select search terms from a pre-
 4 formatted list that is presented to the user at the terminal 14. More sophisticated use of
 5 voice recognition may include actually speaking letters or numbers, or full search terms,
 6 such as speaking a key word for a key word search, for example.

7 The protocol analyzer 133 provides an output 135 to the constraint collator 136,
 8 and the constraint collator 136 determines the nature of the request, block 254. If the
 9 request 114 is a refresh request (i.e., a command to initiate the refresh function), the
 10 constraint collator 136 sends a reset command 131 to the database qualifier 160. The
 11 updated request 115 (possibly with a new constraint) is then sent to the query analyzer
 12 150 for further processing, including analyzing the database 12, retrieving field
 13 descriptors, and formatting, block 256. The result of the data field descriptor retrieval
 14 and formatting are shown as an available data fields result list, block 258, and is returned
 15 to the terminal 14, block 260.

16 In block 254, if the request 114 is not a refresh request, the constraint collator 136
 17 provides the updated request 115 (which may be an initial request, or a subsequent
 18 request) to the query generator 150, block 264. The constraint collator 136 compares the
 19 request 114 against information stored in the status control 140. In particular, the
 20 constraint collator 136 sends the request status control signal 118 to the status control 140
 21 and receives the request status response 119. The constraint collator 136 then compares
 22 the request status response 119 to constraint information provided with the request 114 to
 23 determine if the constraint status should be updated (e.g., because the request 114
 24 includes a new constraint). If the constraint status should be updated, the constraint
 25 collator 136 calls create new constraint subroutine 270, and creates new constraints.

26 The create new constraints subroutine 270 is shown as a flowchart in Figure 17.
 27 The subroutine starts at 272. In block 274, the constraint collator 136 determines if the
 28 request is for a sort-on-the-fly operation. If sort-on-the-fly has been selected, field
 29 assessor 162 prepares a new set of data fields, block 280. The new set of data fields are
 30 then sent to the query generator 150, block 284, and the subroutine 270 ends, block 286.

31 In block 274, if sort-on-the-fly was not selected, the request analyzer 130 may
 32 receive a key word constraint, block 276. The query generator 150 will then generate an
 33 input window in which the user may enter a desired key word, block 282. Alternatively,
 34 the query generator 150 may prompt the user to enter a key word using voice recognition

1 techniques, or any other way of entering data. The process then moves to block 284. In
2 block 276, if a key word search option was not selected, the constraint collator 136 enters
3 the new constraint to the existing list of constraints, block 278. The process then moves
4 to block 284.

5 Returning to Figure 16, the constraint collator 136 next updates the status control
6 140, block 290. In block 292, using the updated constraints, the query generator 150
7 generates a next query of the database 12, block 292. The database driver 170 then
8 extracts the result list from the database 12, according to the latest query, block 294. In
9 block 296, the truncator 152 determines if the result list may be displayed at the terminal
10 14. If the result list cannot be displayed, the process moves to block 298, and a truncation
11 routine is executed. The process then returns to block 294. If the result list in block 296
12 is small enough, the result list is provided by the dispatcher 154 to the terminal 14, block
13 258.

14 As noted above, the request analyzer 130 determines the nature of the request,
15 including any special commands. A special command may include a command to
16 conduct a search-on-the-fly. Alternatively, the search engine 125 may adopt a search-on-
17 the-fly mechanism as a default value. The search engine 125 also may incorporate other
18 special search commands, such as a Boolean search, for example.

19 Figures 18 - 20 are flowcharts illustrating alternate truncation subroutines 298. In
20 Figure 18, the subroutine 298 adjusts a size of a data field by decrementing a parameter
21 TP related to entries in a selected data field. For example, if the data field comprises a list
22 of U.S. cities by name, the parameter TP may be the number of alphabetical characters in
23 a name. The results of such a truncation is shown in the example of Figure 4. The
24 subroutine 298 starts at block 301. In block 303, the parameter TP is set to equal a size of
25 the data field being searched. The truncator 152 then determines the list of records sized
26 by the parameter TP, block 305. In block 307, the truncator 152 determines if the result
27 list can be displayed at the terminal 14. If the result list cannot be displayed at the
28 terminal 14, the truncator 152 decrements the parameter TP, block 309. Processing then
29 returns to block 305, and the truncator 152 gets a reduced result list using the truncated
30 parameter TP. If the result list can be displayed at the terminal 14, the process moves to
31 block 311 and the subroutine 298 ends.

32 Figure 19 is a flowchart illustrating an alternate truncation routine 298'. The
33 process starts in block 313. In block 315, the truncator 152 sets the parameter TP to a
34 size of the data field being searched. In block 317, the truncator 152 determines the list

1 of records sized by the parameter TP. In block 319, the truncator 152 determines if the
2 result list can be displayed at the terminal 14. If the result list cannot be displayed, the
3 truncator 152 adjusts the size of the data field by dividing the parameter TP by a set
4 amount, for example, by dividing the parameter TP by two, block 321. Processing then
5 returns to block 317, and repeats. If the result list can be displayed at the terminal 14, the
6 process moves to block 323 and the subroutine 298' ends.

7 Figure 20 shows yet another alternative truncation subroutine 298". The process
8 starts in block 325. In block 327, the truncator 152 sets the parameter TP to equal the size
9 of the data field being searched. In block 329, the truncator 152 determines the list of
10 records sized by the parameter TP. The truncator 152 then determines if the result list can
11 be displayed at the terminal 14, block 331. If the result list cannot be displayed at the
12 terminal 14, the truncator 152 determines if the parameter TP is less than ten, block 333.
13 If the parameter TP is not less than ten, the truncator 152 adjusts the parameter TP by
14 multiplying the parameter TP by a number less than one, block 337. In an embodiment,
15 the number may be 3/4. The process then returns to block 329 and repeats. In block 333,
16 if the value of the parameter TP is less than ten, the truncator 152 decrements the
17 parameter TP by one, block 335. Processing then returns to block 329 and repeats. In
18 block 331, if the list can be displayed at the terminal 14, the process moves to block 339
19 and the subroutine 298"ends.

20 The examples illustrated in Figures 18 - 20 are but a few examples of the
21 truncations subroutine. One of ordinary skill in the art could conceive of other methods
22 to adjust the field size. In addition to using a truncation subroutine, the user may specify
23 a limit for the field size.

24 As noted above, the search engine 125 may be used for multiple searches and may
25 be used to search multiple databases, including databases with different schemas. The
26 results of individual searches, including the control data provided in the status control
27 140, are saved. The search engine 125 may then be used to further sort (search), or
28 otherwise operate on, the results of these multiple searches. In an embodiment, the search
29 engine 125 may perform a Boolean AND operation on two search results. The result of
30 the Boolean AND operation would be a list of records, or entries, that are common to the
31 two search results. Figure 21 illustrates such a Boolean AND operation.

32 In Figure 21, a GUI 400 displays local database selections 410, including a
33 database of recordings (compact discs - CDs) 412 and a database of contacts 414. The
34 databases 412 and 414 may be shown by text descriptions and an appropriate icon, for

1 example. The database selections in this example are resident on a user's terminal, such
2 as the terminal 14 shown in Figure 1. Also displayed on the GUI 400 is a remote
3 database selection 420 that represents databases, such as the databases 13 and 15 shown
4 in Figure 1, that are located remotely from the terminal 14. In the example shown in
5 Figure 21, the remote database selection 420 includes a database 422 for online record
6 sales, which is represented by an icon (a CD) and a text title of the online retailer. The
7 remote databases shown in the remote database selection 420 may include those databases
8 for which the user has already established a link. In the example shown, the user may
9 already have entered an Internet address for the online retailer. In addition to any
10 returned web pages from the online retailer, the terminal 14 may then display a
11 representation of the database 422.

12 Continuing with the example, the user may use the search engine 125 to conduct a
13 search-on-the-fly of the recordings database 412 and the Virgin Records™ database 422.
14 The user may search both databases 412 and 422 for titles of recordings that are classified
15 as "blues." The search engine 125 may return search results 416 and 424 for searches of
16 both databases 412 and 422, respectively. The search results 416 and 424 may be
17 displayed in a window section 430 of the GUI 400. The results 416 and 424 may also be
18 represented by CD icons, such as the icons 432 and 434. The search results 416 and 424
19 may be stored as lists in one or more temporary databases, as represented by the windows
20 417 and 427. The search results 416 and 424 may also be stored in a scratch pad database
21 418. At this point, the user may wish to determine which recordings from the list 424 are
22 contained in the list 416. The search engine may support this function by performing a
23 Boolean AND operation of the lists 416 and 424. The results of the Boolean AND
24 operation are represented by the icon 436 displayed in the window 430. To execute the
25 Boolean AND operation, the user may simply drag the icon 432 over the icon 434, and
26 then select AND from a pop-up menu 438 that appears when the icons 432 and 434
27 intersect. Other techniques to execute the Boolean AND (or another Boolean function)
28 may include typing in a command in a window, using voice recognition techniques, and
29 other methods. In addition, other Boolean functions may be used.

30 The result represented by the icon 436 of the Boolean AND operation may then be
31 stored in a database at the terminal 14, such as in the scratch pad database 418 or may be
32 stored at another location. The result may then be subjected to further search-on-the-fly
33 operations.

Also shown in Figure 21 is an online-purchase module 435 that may be used to consummate purchase of a product referenced in an online database such as the database 422. To initiate such a purchase, the user may drag an iconic or text representation of a desired product listed in the search result 424 over an icon 436 in the online-purchase module 435. This drag-and-drop overlaying these icon may initiate and complete the online purchase for the desired product.

Use of the search engine 125 may be facilitated by one or more GUIs that are displayed on the terminal 14. Figures 22 - 26 are examples of such GUIs. In Figure 22, a GUI 450 includes a display section 452 and one or more database sections such as local database section 470 and remote database section 460. The local database section 470 includes databases local to the terminal 14. In the example shown, the local databases include a patients database 472, a general contacts database 474, a pharmacy database 476, a medicines database 478 and a scratch pad database 480. The remote databases include an Amazon.com database 462, an online record retailer database 464, a Physician's Desk Reference database 466 and an American Medical Association (AMA) online database 468. The remote and local databases may be represented by a text title and an icon, both contained in a small window as shown. A user may access one of the remote or local databases by moving a cursor over the desired window and then selecting the database. In the example shown, the local medicines database 478 has been selected, and a list 490 of data fields in the medicines database 478 is displayed in the display section 452. Also included on the display section 452 is a keyword button 492 that may be used to initiate a key word search of the medicines database 478.

Figure 23 shows the GUI 450 with a user selection of a category data field from the list 490. The category data field is indicated as selected by an arrow adjacent to the data field name. When the category data field is selected, a category list 494 is displayed on display section 452. The category list 494 includes four entries, as shown.

The user may continue to search the medicines database 478 using key word techniques and search-on-the-fly techniques. Figure 24 shows the GUI 450 with results of several search cycles displayed.

Figure 25 illustrates a search of the PDR database 466. Such a search may be initiated by dragging a cursor to the window having the PDR 466 symbol (text or icon), and then operating a “select” button. Figure 26 shows a search of the Amazon database 462. This search may also be initiated by a “drag-and-drop” operation.

1 The SOTF search engine 125 may accommodate merging of one or more sets of
2 search results. The multiple search results may be derived from a common database, or
3 from more than one database. A search using the search engine 125 may be controlled
4 through a user interface by one or more icons that can represent (1) filters or (2) the
5 images of filters. Thus, the icon may represent spatial or temporal attributes, or sets of
6 objects or procedures. Merging the icons thus has two interpretations corresponding to
7 (1) and (2): either filters are added ("apply every filter in every icon to every image to
8 which it can be applied"), or image sets are added. In an alternative embodiment, the
9 addition (union or join) operator may be any other relational operator, e.g. divide,
10 difference.

11 Use of the merge function may be explained by reference to Boolean lattices. A
12 collection of entities can have attributes A or B or both. If $\{A\}$ is the set of all A entities
13 and $\{B\}$ is the set of all B entities; the set whose elements all possess both attributes A
14 and B may now be written $\{A \text{ and } B\}$, and the set whose elements all possess either
15 attribute A or attribute B or both may be written $\{A \text{ or } B\}$. The elements of $\{A \text{ and } B\}$
16 can be considered to possess a new, less inclusive or specific attribute C, and the elements
17 of $\{A \text{ or } B\}$ to possess a new, more inclusive or general attribute D. In a lattice, the
18 nodes are attributes; the most inclusive attribute (in this case D) is always at the top and is
19 called the join of those attributes (nodes) immediately below it, and the most exclusive
20 attribute (in this case C) is always at the bottom and is called the meet of those attributes
21 (nodes) immediately above it. In other words, the OR operation yields the join of two
22 attributes, while the AND operation yields their meet. Thus, the OR operator is upward or
23 inductive (yielding the more inclusive join of the operands), while the AND operation is
24 downward or deductive (yielding the more exclusive meet of the operands). The nodal
25 attributes of such a lattice are analogous to filters; but since a principle called CF duality
26 states that attributes and sets are to some extent interchangeable because every attribute
27 characterizes a set and every set is characterized by an attribute, these attributes are
28 logically equivalent to the sets they characterize.

In an example optical context, the downward AND operator corresponds to stacking colored filters, while the upward OR operator corresponds to mixing colored paints or filters. In color optics, stacking and unstacking colored lenses is called a subtractive process, while mixing or unmixing paints is called an additive process. Unfortunately, while combining or “adding” filters is subtractive with respect to the sets they characterize, it is additive with respect to the filters themselves, and adding sets is

1 subtractive with respect to the filters. So it is better to refer to operations among
 2 attributes (filters, lenses, etc.) as “filtrative” or “infonegative, and to those among sets
 3 (paints, lights, etc.) as “constructive” or “infopositive”. CF duality can now be rephrased
 4 as follows: every infonegative entity (attribute) descriptively characterizes an associated
 5 infopositive entity (set/object), and every infopositive entity instantiates or is
 6 descriptively characterized by an associated infonegative entity.

7 The search engine 125 may include iconization (iconic representation) of an
 8 algebra or calculus of relations defined on Boolean lattices. This representation begins
 9 with a set of primitive icons extracted from base tables and defines new icons (derived
 10 tables, virtual databases) by means of simple user-executed operations. The icons can be
 11 effortlessly translated into lists of data corresponding to the icons, and it is these lists that
 12 comprise the real substance of any search procedure.

13 When search chains are branched into to chains A and B, the filters subsequently
 14 applied to each chain can be the same or different, and merging can signify any of two or
 15 more Boolean relationships (relational operations) defined on a relational database.
 16 Specifically, when chains merge, sets of filters can be added or intersected. Since filters
 17 are constraints, adding them amounts to intersecting their images, while adding their
 18 images amounts to intersecting the filters (infopositive-infonegative distinction).
 19 Equivalently, one may consider positive and negative filters effecting deduction and
 20 induction respectively; the filters are descriptive, while the images are substantive. The
 21 extent to which the images of filters can intersect depends on the commonality
 22 (predicative non-exclusivity) of domains. Icon algebras (of iconic operators) are “object-
 23 oriented” on the GUI level; they are UI extensions of the innate object-orientation of
 24 relational databases themselves, wherein the objects are records, attributes, tables, virtual
 25 databases and so on, and the operations are those of any relational algebra.

26 The looping and merging of search chains is to some extent algebraic. First, since
 27 actual topology is being changed, such transformations do not directly form a topological
 28 homeomorphisin group; the algebra remains Boolean, and the “homeomorphism” is
 29 defined on the operator graph of the Boolean algebra (of which the initial search tree is
 30 generally only a subspace). Icons representing sets of nested predicates are “Boolean
 31 objects”; when decision chains converge or diverge, objects merge or split, and these
 32 objects represent (combinatorially) unique search paths. Thus, operations among paths
 33 can be reduced to operations among objects; e.g., regress-diverge is just an object-
 34 splitting operation. Continuous looping applies “inverse deductive filters” to achieve

1 induction by descriptive intersection of filter constraints, permitting the retrograde
 2 convergence of paths to identical ancestral objects (inductive merging of objects), while
 3 inductive looping is just direct regression to an ancestral object preparatory to splitting it
 4 and thus effecting divergence of paths (deductive splitting of objects). Deductive
 5 convergence of paths is "natural" if iconic image sets intersect and "forced" if not; if
 6 natural, then there has been non-exclusivity of subobjects, and paths are not unique (even
 7 though identical filters can apply to divergent paths without impairing uniqueness). So all
 8 deductive merging is forced, and this entails a decision regarding which filters are to be
 9 conserved and which discarded. Any such operation will effectively "rewrite the paths",
 10 and doing this optimally is NP-complete.

11 More specifically, icons are subject to CF duality. The merge control thus has a
 12 "switch" toggling between "Qualities / Objects". When the switch is in the "qualities"
 13 position, merging icons performs a qualified deductive conjunction of filters and yields a
 14 set intersect; when it is in the "objects" position, merging the icons performs a disjunction
 15 of filters and an inductive union of sets, yielding a more general attribute (the general
 16 qualities created by the object-merge operation will be produced by sets of filters applied
 17 disjunctively). The search engine 125 is therefore capable of inductive and deductive
 18 information processing. A quality-merge in which filters do not cross the line between
 19 composite icons equates to an object merge; the set thus selected is characterized by a
 20 more general quality which amounts to the descriptive (filtrative) union. There is also a
 21 modified quality-merge in which filters in either icon applicable to both iconized sets are
 22 applied to both, thus crossing the line between icons. In this case, a true merging of paths
 23 occurs, as opposed to path icons. The search engine 125 allows users to choose which
 24 filters are to cross the inter-icon line and which are not, resulting in complex Boolean
 25 expressions and the sets they characterize (determining consistency of complex
 26 expressions can amount to LSAT; sets of inconsistent expressions will simply yield a null
 27 return.

28 Icons may reside in the first menu box to appear, being transferred from menu to
 29 menu as the path is generated and filters are accumulated. When a direct regress occurs,
 30 the path is regarded as "complete" and is stored in a holding module. Prior to the
 31 merging operation, the quality/object switch is set; and icon subfilters or subsets
 32 individually displayed. A "lattice navigator" will keep track of position and equivalence,
 33 folding the search graph in case a node of the original tree is inductively encountered in
 34 the course of an object-merge; otherwise, the icon remains in "internodal space" (which is

1 to be regarded as a virtual space realized only in the event that the search tree is
2 nondisjunctive in its nodes and therefore incomplete with respect to the semantic net
3 generated by the tree).

4 Figure 27 is a flow chart illustrating an alternative operation 600 of the query
5 generator 150 of Figure 6. In the illustrated operation, the query generator 150 is adapted
6 to receive multiple selections of items within a same menu function and within a same
7 merge function. To provide this functionality of the query generator 150, the request
8 analyzer 130 (see Figure 5) may be adapted to receive a collection of user choices.

9 The operation 600 begins in block 601. In block 603, the request analyzer 130
10 receives constraints collected from the constraint collator 136, and the updated request
11 115, which may be an initial request or a subsequent request, is provided to the query
12 generator 150. In block 605, the query generator 150 determines if the constraints (the
13 request 115) are in the same merge group. If the query generator 150 determines that the
14 request 115 is in the same merge group, the process moves to block 607 and the query
15 generator 150 generates the query with a Boolean AND. If the request is not in the same
16 merge group, the query generator 150 generates the query with a Boolean OR, block 609.

17 In block 611, the items selected within the same unit are Or'ed and the default
18 truncator may be used depending on the size of the returned items. In block 613, the
19 generated query is executed. In block 615, the number of records to be displayed is
20 checked. If the number is within a specified limit, the process moves to block 617 and
21 the search results are returned for display. The operation 600 then ends, block 625. In
22 block 625, if the number of records to be displayed is too large, the process moves to
23 block 621, and a truncation routine is executed.

24 The truncation routine may be any of the previously-described truncation routines
25 illustrated in Figures 18-20. Figure 28a illustrates an alternate truncation routine 630.
26 The routine 630 begins in block 631 with the truncator 152 receiving the request 115. In
27 block 633, the truncation is set to the size of the field being viewed on the GUI, and sets
28 the False Flag. The query is then run against the database using the selected truncator,
29 block 635. In block 635, the truncator 152 determines if the number of records that
30 would be retrieved from the database can be displayed on the existing GUI. If the records
31 can be displayed, the process moves to block 639, and the truncator 152 determines if the
32 Flag is set False. If the Flag is set False, the process moves to block 653 and the records
33 are returned (displayed on the GUI). The process then ends, block 655. In block 637, if
34 the number of records exceeds the display size of the GUI, the status of the Flag is

1 checked as False. If false, the truncator is set to 1, and the flag is set to true, block 647,
2 and the process returns to block 635. If in block 637. If the flag is not set false, the
3 process moves to block 651, and saved records are retrieved. The retrieved records are
4 then displayed, block 653.

5 In block 639, if the Flag is not set to false, the retrieved records are saved, and the
6 truncator 152 is incremented. The process then returns to block 635.

7 Figure 28b illustrates another alternative truncation routine 700. In block 701, the
8 truncator 152 receives the constraints, the view by field and the maximum of number of
9 display items (MNDI). In block 702, the truncation is set to zero (no truncation), and the
10 Flag is set to True. Next, the query is generated in block 702. In block 703a, query
11 generator receives the constraints, the view by field, and the truncator as parameters, and
12 the query generator returns the query. The query is then run against the database, and the
13 counter is set to zero, block 704. In block 705, the truncator 152 fetches the next record
14 and increments the counter. If the end of file is reached, block 706, and the truncation
15 equals zero, block 710, the truncator 152 sends the list of fields to the client, block 712.
16 However, if the truncation is not zero, block 710, the truncator 152 is incremented, block
17 709, and the process returns to block 703. On the other hand, if the end of file is not
18 reached, block 706, and the counter is smaller than MNDI, block 707, the process goes
19 back to block 705, in which the truncator 152 fetches the next record and increments the
20 counter. However, if the counter is larger than MNDI, block 707, and the saved list of
21 fields exist, block 708, the truncator sends the list to the client, block 712. Conversely, if
22 the saved list of fields do not exist, block 708, the truncator 152 is incremented, block
23 709, and the process goes back to block 703 again.

24 Table 1 illustrates an example of the alternate truncation routine 700. This routine
25 begins by attempting not to truncate the records.

1 Table 1

| Records | | 1 st Round | | 2 nd Round | | 3 rd Round | |
|---------|-------------|-----------------------|-------------|-----------------------|---|-----------------------|----|
| 1 | Armandia | 1 | Armandia | 1 | A | 1 | AR |
| 2 | Armonk | 2 | Armonk | 2 | N | 2 | NE |
| 3 | Armonk | 3 | New Orleans | 3 | R | 3 | RI |
| 4 | New Orleans | 4 | New York | | | 4 | RO |
| 5 | New Orleans | | | | | | |
| 6 | New York | | | | | | |
| 7 | New York | | | | | | |
| 8 | New York | | | | | | |
| 9 | Riverdale | | | | | | |
| 10 | Riverdale | | | | | | |
| 11 | Riverdale | | | | | | |
| 12 | Rockfort | | | | | | |

2 In this example, the maximum number (n) of displayable results is three, and the
3 database contains twelve instances of six different cities. First, the database is queried for
4 the full city field with no truncation, and records are fetched. Records are fetched until
5 four (n+1) records are fetched from the database. Since the number of different cities (4)
6 is greater than n, fetching is halted and the process moves to truncation. Then the
7 database is queried for only the first letter of the cities (truncation is incremented so that it
8 equals one). For this query the database manager may simply review its index. The
9 compiled list from the query is saved as "A", "N", and "R". Next, the database is queried
10 for the first two letters of the city field (truncation is incremented so that it equals two).
11 Again, the database manager may simply review its index to locate this information about
12 the data field. This query for two letters or characters is continued until the number of
13 two letter combinations exceeds n. When the number of different combinations (4) is
14 again greater than n, the routine halts and nothing is saved. The system now returns to
15 the previous saved list. Therefore, the saved list ("A", "N", and "R") is returned to the
16 client for display or process.

17 Figures 29 - 38 illustrate graphical user interfaces and search on the fly results
18 using the search engine 125 with a merge function. In Figure 29, a search of a patent
19 database has been executed to search for patents by primary examiner. The Primary
20 Examiner results table lists the arabic numerals 0 - 7 and the letters A-Z, indicating that
21 the database contains names of primary examiners beginning with these numerals/letters.
22 To quickly narrow the search, the user selects the letter O, and results are returned listing
23 last and first names all primary examiners whose last name begins with O. As can be
24 seen by the returned results, the database lists several primary examiner instances of
25 O'Dea. This could indicate an error in the database. The search engine 125 allows these

1 errors to be detected and corrected. The correction may be made by selecting the
2 incorrect instances, right-clicking the correct instance, and then choosing a 'correct all
3 other's based on this instance" function.

4 Figure 30 shows how multiple-select capabilities of the search engine 125 may be
5 used to enhance a search. In the illustrated example, the user searches for 3M Company.
6 Different versions of the company name are then displayed with the returned results. In
7 this way, the user may select the different versions of the company that the user wants to
8 use for the search. The pop-up pane shows a current status control for the GUI.

9 Figure 31 shows the results of subsequent menus showing the aggregation, or
10 merge, of two previous constraints, "3m" and 3-M." Figure 32 shows a merge execution.
11 The user first selects the '3-M" and the "3M" company names using the check boxes in
12 the previous menu. The user then selects the merge option, placing the menu on hold,
13 and going to the "M", "MI", "MIN" and "MINNESOTA M" menus. The merge option is
14 then selected on the menu and the merged menu is displayed showing the merge of
15 searches between "3M" and "Minnesota Mining and Manufacturing Co." Figures 32 - 36
16 show other search engine 125 features including data mining and database correction.

17 Figures 37 - 39 show the results of a full text search of a patent database using the
18 keyword "encryption" and searching on all fields. The initial search results are truncated
19 to display by first letter/numeral of the patent title. From this intermediate search result
20 menu, the user selects all patents whose title begins with the letter "E", and a subsequent
21 search result menu is displayed listing partial titles of all such patents. From the next
22 intermediate list, the user selects the patent whose title begins "Electronic copy protection
23 mechanis." (see Figure 38) The search engine 125 then returns this specific patent, the
24 first page of which is shown in Figure 39. The displayed patent includes the keyword
25 "encryption" highlighted wherever it occurs. The display also indicates the number of
26 instances of the keyword in the patent.

27 Figures 40-49 illustrates additional search results.

28 In the examples shown in Figures 37-49, search results are displayed on a "large-
29 format" screen, such as available with a desktop personal computer. When a user is in a
30 mobile environment (e.g., on foot, in a car) the user may still be able to access the search-
31 on-the-fly search engine and have search results returned to a mobile display device such
32 as a cellular telephone or a personal data assistant.

33 Figure 50 illustrates a standard cellular telephone 800 that may use the search-on-
34 the-fly search engine 125. The cellular telephone 800 includes a display 801, a keypad

1 802, and other controls 803 that may be used to navigate one or more data buses using the
2 search-on-the-fly search engine 125.

3 Figure 51 illustrates a personal data assistant (PDA) 800 that may use the search-
4 on-the-fly search engine 125. The PDA 800 includes a display area 811 and an input area
5 812.

6 Figures 52a – 52l illustrate a search sequence using the cellular telephone 800
7 configured to use the search-on-the-fly search engine 125. In the example illustrated, the
8 U.S. Patent and Trademark Office patent database is selected. Using the cellular
9 telephone 800, the user conducts a search of the U.S. Patent and Trademark Office
10 database using a series of filters. Each time a filter is applied, a search result may be
11 returned and displayed on the display 801. Using the controls 802, the user may add or
12 subtract filters. The display 801 shows the accumulative result of the filtering process.
13 When the data to be returned is too large to fit the display 801, the returned data may be
14 truncated as illustrated in Figures 52f-52k.

15 Figure 53 illustrates a general purpose personal computer system 850 that may be
16 used for search-on-the-fly of a plurality of databases. The system 850 includes a
17 processor section 851, a display and a control section coupled to the processor section
18 851, and a computer readable medium 855, which may be read by components of the
19 processor section 851. The computer readable medium 855 may include the software
20 routine required to implement the search-on-the-fly with merge function method.

21 In specific embodiments, the search engine 125 is implemented as a program
22 executed on a general purpose computer, such as a personal computer. The search engine
23 may also be implemented as a routine attached to a database structure. In addition, the
24 search engine may be implemented on any processor capable of executing the routines of
25 the program. In alternative embodiments, the search engine 125 may be implemented as
26 a single special purpose integrated circuit (e.g., ASIC) having a main or central processor
27 section for overall, system level control, and separate circuits dedicated to performing
28 various different specific functions, computations and other processes under control of the
29 central processor section. Those of ordinary skill in the art will appreciate that the search
30 engine 125 may also be implemented using a plurality of separated dedicated or
31 programmable integrated circuits, or other electronic circuits or devices (e.g., hardwired
32 electronic or logic circuits such as discrete elements circuits, or programmable logic
33 devices, such as PLDs, PLAs, or PALs). In general, any device or assembly of devices

1 on which a finite state machine capable of implementing flowcharts similar to the
2 flowcharts of Figures 16 – 20, 27 and 28 can be used to implement the search engine 125.

3 While using search on the fly has been described in detail for an end result of
4 printing, viewing or displaying data, search on the fly can be useful for other purposes.
5 Search on the fly does not require obtaining the underlying data in the database or the
6 display of the underlying data to be useful. Search on the fly can be used for gathering
7 information or characteristics about data in a database with or without downloading the
8 data itself. This gathered information about the data can be used to analyze the data,
9 sorting, correct or clean data, verifications and confirmations. For example, search on the
10 fly can be used to determine whether there is existing data in a database within certain
11 ranges or parameters (date ranges, numerical, alphanumerical and other characteristics).
12 If there is data within certain parameters, the number of datapoints within those
13 parameters can also be determined. This information about the data can be gathered
14 using search on the fly with queries to the database manager (which may simply need to
15 query its index and not access the data itself). Another example is correcting data. Data
16 may need to be corrected or cleaned for various reasons including spelling errors. Search
17 on the fly can locate these errors without necessarily accessing and downloading the data
18 itself. Certain combinations of characters or truncations will be obvious spelling errors.
19 Also, data that is out of range can be located and corrected or eliminated from the
20 database using search on the fly. Another example is data from one database can be
21 confirmed or verified against data in a second database using search on the fly. Those
22 skilled in the art will find many uses and specific applications for search on the fly.

23 The terms and descriptions used herein are set forth by way of illustration only
24 and are not meant as limitations. Those skilled in the art will recognize that many
25 variations are possible within the spirit and scope of the invention as defined in the
26 following claims, and there equivalents, in which all terms are to be understood in their
27 broadest possible sense unless otherwise indicated.